

By Srinivas Thodati
Brian Johnson

MINIMIZING NETWORK COMPLEXITY IN VMWARE vSPHERE 4 ENVIRONMENTS WITH INTEL 10 GIGABIT ETHERNET CONNECTIONS FOR DELL POWEREDGE SERVERS

Today's 10 Gigabit Ethernet server adapters help to greatly simplify networking in VMware® vSphere™ 4 virtualized environments. IT organizations can take advantage of advances in both Intel® Ethernet technology and vSphere software to create a flexible, highly efficient networking environment without compromising areas such as security and traffic segmentation.

The success of virtualization, while substantial, has been limited by the complexity that has arisen from networking with Gigabit Ethernet (GbE) technology. Networking virtualized hosts with GbE is typically based on relatively large numbers of dedicated physical connections, which are used to segregate different types of traffic because of networking infrastructure and bandwidth limitations from 1 Gbps connectivity. Organizations pay for this complexity in increased management and equipment costs, as well as increased power usage.

As 10 Gigabit Ethernet (10GbE) technology enters the mainstream, IT organizations are considering how they can use it to help simplify their virtualized environments. Consolidating multiple GbE traffic flows onto a reduced number of 10GbE connections helps reduce complexity and cost. However, many IT administrators have expressed concerns that quality of service (QoS), security, or performance may suffer as a result of this traffic convergence. These concerns have prevented some organizations from taking full advantage of 10GbE technology.

Cumulative experience among Dell, Intel, and VMware together with advances in Dell™ PowerEdge™ servers, Intel Ethernet 10GbE server adapters, and VMware vSphere 4 virtualization software can alleviate

these concerns and help organizations realize further efficiency and cost savings.¹

TRANSITIONING TO A VIRTUALIZATION MIND-SET

The common practice when deploying virtualized hosts in a physical environment has been to segregate network functions onto dedicated GbE ports, adding additional ports as demand or bandwidth increases. These ports are often installed in pairs to provide network failover, doubling the number of ports required per host. As a result, the number of network ports has a tendency to become bloated, leading to excessive complexity and associated high costs of hardware and power.

This approach is largely a carryover from physical configurations. Pre-virtualization best practices guided initial implementations of virtualization technology, and physical server network paradigms were often extended to virtual infrastructures. This extension led to the use of separate physical connections for segmenting traffic and providing the required bandwidth.

In addition, previous versions of the VMware platform required dedicated connections for virtual machines (VMs) and for each of multiple traffic types, such as VM traffic, service console connections, IP

¹To learn more about optimizing performance in virtualized environments with 10GbE connectivity, see "Intel Ethernet Server Adapters Maximize Throughput Performance for iSCSI Connectivity," by Sunil Ahluwalia and Gary Gumanow, in *Dell Power Solutions*, 2010 Issue 1, dell.com/downloads/global/power/ps1q10-20100322-Intel.pdf.

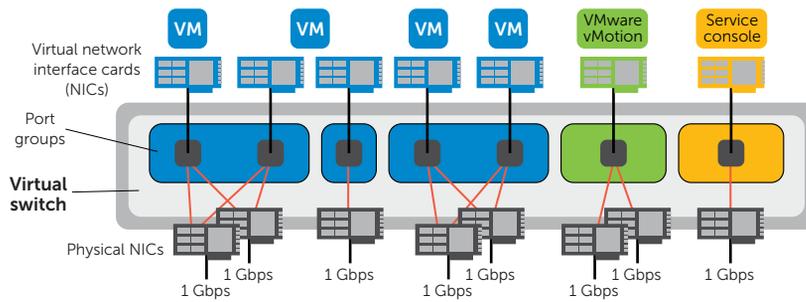


Figure 1. Virtual switch with multiple physical GbE server adapters

storage, and VMware vMotion™ technology. Security procedures also led network administrators to physically segregate traffic onto separate ports, because the IEEE 802.1Q standard for trunking to the host was not implemented as a standard practice and was limited to switch-to-switch connections.

Segregating a 10GbE port into multiple dedicated bandwidth channels, however, can significantly reduce the advantages of moving to a 10GbE environment. For example, dividing a connection into four channels with dedicated bandwidths of 4 Gbps, 2 Gbps, 1 Gbps, and 1 Gbps means that no single connection can use more than those limits, because the channels cannot share unused bandwidth. Unlike using one open connection in which unused bandwidth can be used by other traffic types, segregating the bandwidth of a 10GbE connection does not provide significantly more headroom than using multiple GbE connections, and leads to a significant increase in management overhead. If a maximum of 4 Gbps of throughput is all an organization needs on the largest connection, then a single open 10GbE connection can handle any of the traffic with plenty of bandwidth to spare, especially in a typical implementation using 10GbE ports in a redundant active/active configuration.

COMPARING GbE AND 10GbE TOPOLOGIES

Using GbE, host servers may need as many as eight or more network ports to satisfy the requirements of virtualization (see Figure 1). In this topology, several

port groups are configured to support the various networking functions and application groupings. In turn, each port group is supplied with one or more physical connections. Virtual LAN (VLAN) tags may be implemented on these port groups as well. This GbE topology raises several issues, including the following:

- **Inefficiency:** The large number of physical ports and server adapters contribute to increased administrative overhead and power consumption.
- **Difficult network management:** The increase to 8–12 network ports on each server can increase the likelihood of misconfiguration.
- **Numerous points of potential failure:** Multiple physical devices and cable connections contribute to the likelihood of hardware failures.
- **Bandwidth limitations:** Static allocation and physical reconnections are required to add bandwidth to the network.

Moving from multiple GbE connections to fewer 10GbE connections enables

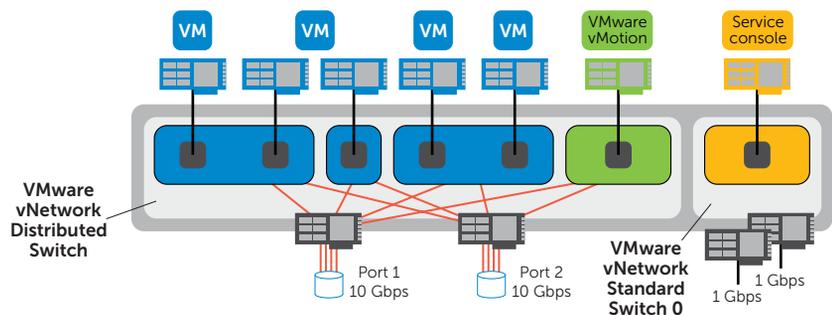


Figure 2. VMware vNetwork Distributed Switch with 10GbE server adapters for network traffic and GbE server adapters for service console traffic

a flexible, scalable network infrastructure that helps reduce complexity and management overhead while providing high availability and redundancy. The 10GbE installation shown in Figure 2 is analogous to the GbE topology shown in Figure 1 but uses 10GbE connectivity for most ports (limiting GbE to the service console) and capitalizes on the VMware vNetwork Distributed Switch feature in VMware vSphere 4. This feature provides the same basic functions as standard virtual switches, but the switches exist across two or more clustered VMware ESX or ESXi hosts. (For an overview of performance-enhancing technologies for flexible infrastructures, see the “Enabling Near-Native 10 Gigabit Ethernet Performance” sidebar in this article.)

ADDRESSING SECURITY, SEGMENTATION, AND BANDWIDTH CONCERNS

Even though 10GbE networking enables the consolidation of multiple functions onto a single network connection, the practice of using large numbers of GbE connections has persisted because of several administrator concerns—including segregating traffic for security, meeting QoS requirements, and providing dedicated bandwidth for critical networking functions. When GbE server connections are consolidated onto 10GbE connections, isolating connections in the absence of dedicated physical connections is still necessary. This requirement reflects the need for security between different types of traffic. Isolating connections is also

ENABLING NEAR-NATIVE 10 GIGABIT ETHERNET PERFORMANCE

Realizing the performance benefits of 10 Gigabit Ethernet (10GbE) in virtualized environments requires minimizing software overhead. Several technologies can enhance networking performance in these environments:

- **Intel Virtual Machine Device Queues (Intel VMDq):** Offloads data-packet sorting from a virtual switch onto a physical network adapter, helping to reduce processor overhead and enhance overall efficiency*
- **Single Root I/O Virtualization (SR-IOV):** Allows administrators to partition a physical port into multiple virtual I/O ports, and use each as a single dedicated port directly assigned to a virtual machine
- **Intel Virtualization Technology for Directed I/O (Intel VT-d):** Enables an I/O resource to be assigned to a specific virtual machine, giving its OS direct access to that resource**
- **Data Center Bridging (DCB):** Provides enhanced traffic prioritization over a single interface and, in conjunction with 10GbE, can help cost-effectively meet latency requirements***

*To learn more about Intel VMDq, see "Intel Ethernet Server Adapters Maximize Throughput Performance for iSCSI Connectivity," by Sunil Ahluwalia and Gary Gumanow, in *Dell Power Solutions*, 2010 Issue 1, dell.com/downloads/global/power/ps1q10-20100322-Intel.pdf.

**To learn more about Intel VT-d, see "Building the Efficient Enterprise with Dell, Intel, and VMware vSphere 4," by Andrew Gilman, Mike Monthei, and Andrew I. Fields, in *Dell Power Solutions*, September 2009, dell.com/downloads/global/power/ps3q09-20090301-VMware.pdf.

***To learn more about DCB, see "Mixing Gigabit Ethernet and 10 Gigabit Ethernet in a Dedicated SAN Infrastructure," by Tony Ansley, in *Dell Power Solutions*, September 2009, dell.com/downloads/global/power/ps3q09-20090416-Ansley.pdf; and "10 Gigabit Ethernet: Unifying iSCSI and Fibre Channel in a Single Network Fabric," by Achmad Chadran, Gaurav Chawla, and Ujjwal Rajbhandari, in *Dell Power Solutions*, September 2009, dell.com/downloads/global/power/ps3q09-20090392-Chadran.pdf.

necessary to help ensure adequate bandwidth for specific applications within the shared connection.

Employing 10GbE with VMware vSphere 4, as in the infrastructure shown in Figure 2, helps meet these requirements using updated approaches rather than discrete physical connections. VLANs provide the basic security features needed, and VLAN traffic segmentation provides dedicated bandwidth for each type of network traffic. The IEEE 802.1Q standard for VLAN trunking enables administrators to group multiple VLANs onto a single wire, helping reduce complexity. The grouping can be managed as a unit over one physical wire and broken into dedicated networks at the switch level instead of at each host.

USING TRAFFIC SEGMENTATION TO ISOLATE DATA

VLANs allow network traffic segmentation without dedicating physical ports to each segment. This method has the obvious benefit of reducing the number of physical ports needed to isolate traffic types. Dell

and Intel have developed best practices for the 10GbE model to address questions of performance, security, and bandwidth allocation. These best practices are based on extensive experience in the organizations' own IT environments as well as advances in server adapter and virtualization technologies.

Optimum performance through VLAN segmentation

In a virtualized environment, the switch that previously was the first connection out of the server is now part of the hypervisor, and therefore part of the server. In the past, VLAN trunking from physical servers was not typically necessary; with 10GbE and virtualization, however, it can help significantly with port consolidation while still providing traffic segmentation. The key is to segment the traffic to various VLANs and use VLAN trunking to connect the host server to the top-of-rack or end-of-row switch, or to the blade switch for a blade host server.

To help maximize performance, the service console should be on a dedicated

port group with its own dedicated VLAN ID, and should use port 2 in a two-port configuration or a dedicated GbE redundant team. VMware vMotion should also be on its own dedicated port group with its own dedicated VLAN ID, and should also use port 2. IP-based storage traffic—Internet SCSI (iSCSI) and Network File System (NFS)—should be on its own port group in the vNetwork Distributed Switch, using port 2. VM traffic can use one or more VLANs, depending on the level of separation needed between the VMs. In a non-failed state, the VMs should not share the service console, vMotion, or IP-based storage traffic VLAN IDs, and should use port 1 in a two-port configuration.

VLANs for traffic isolation

When VLANs are separated in this manner, logical partitioning isolates individual traffic flows. VMware vSphere can control the effects of individual VMs on the traffic flows of other VMs that share the same physical connection. Demilitarized zones (DMZs) can be configured on different network adapters to isolate internal traffic from external traffic. Administrative traffic and other back-end services are handled by a separate networking stack managed by the VMkernel, providing further isolation from VM traffic even on the same physical connection.²

Enterprise-level reliability

Many IT organizations are concerned about the need to ensure enterprise-level reliability for critical applications and workloads. To enhance reliability, administrators can install 10GbE ports in pairs to support a redundant configuration. If two 10GbE ports are used, a best practice is to run VM traffic primarily on port 1 and all other traffic on port 2. This design uses the bandwidth of both 10GbE ports and can be configured for network failover.

In the event of a hardware failure, the host and management software should migrate all VMs off the host using vMotion to retain redundancy and help ensure

²To learn more about using VLANs to protect data, see "vSphere 4.0 Security Hardening Guide," by VMware, Inc., April 2010, available at communities.vmware.com/docs/DOC-12306.

reliability. To enhance redundancy further, a second option is to move to a configuration of four 10GbE ports, with two primary ports and backup ports on separate adapters.

MEETING THE BANDWIDTH REQUIREMENTS OF CONVERGED DATA STREAMS

Another concern among many administrators is that 10GbE connections can become saturated with one traffic type, negatively affecting other traffic classes. In particular, live migration of VMs using VMware vMotion can be bandwidth intensive. However, reaching saturation on 10GbE is unlikely, given that the converged traffic streams previously used only GbE connections.

Using Dell PowerEdge servers with powerful multi-core processors and Intel 10GbE server adapters can provide the overall bandwidth to support this aggregation. 10GbE adapters provide the bandwidth necessary for multiple traffic types to coexist on a single port. In many cases, QoS requirements can be met simply by the availability of large amounts of bandwidth. The presence of sufficient bandwidth can also increase the speed of live VM migration using vMotion, removing potential bottlenecks for enhanced performance.

In addition to generally available bandwidth, traffic segmentation helps ensure dedicated bandwidth for specific traffic types and QoS. Shared 10GbE server adapters are more advantageous than separate GbE connections for this purpose. The 10GbE connections make more efficient use of available bandwidth than GbE connections for demanding applications and heavy-use portions of the workday.³

GbE limitations with traffic spike management

Although separating traffic flows onto discrete GbE connections is a viable means of providing dedicated bandwidth, doing so has distinct shortcomings. For example, allocating two GbE connections to a VM traffic

port group provides a potential of 2 Gbps of dedicated bandwidth. However, if additional bandwidth is needed for sporadic traffic spikes from that port group, additional server adapters must be added—assuming additional PCI slots are even available. Another shortcoming is that the bandwidth allocated in this example cannot be used by any other traffic, so it simply goes to waste.

10GbE and seamless management of traffic spikes

Providing bandwidth to the port group from a shared 10GbE server adapter allows additional bandwidth to be allocated seamlessly for traffic spikes. Multiple port groups can share the bandwidth headroom provided by the server adapter. The resource can be automatically and dynamically reallocated as various port groups are accessed by their associated VMs. Dedicated bandwidth is not necessary for any one port group as long as the host network connection does not reach saturation.

DEPLOYING 10GbE TO HELP REDUCE COMPLEXITY AND COST

Using Dell PowerEdge servers with Intel 10GbE server adapters in VMware vSphere 4 virtualized environments enables organizations to reduce the complexity and cost of their network infrastructure. These hardware and software technologies are designed to work together to ensure that security and performance requirements can be met without the large numbers of physical server connections required in legacy GbE networks. As a result, administrators can replace methods that depend on physical separation with updated approaches that use logical separation.

The key for IT organizations is to resist trying to mimic a traditional setup of physical servers in a virtualized server environment. Rather than segregating a 10GbE port into multiple dedicated channels, they can take advantage of open 10GbE connections and VMware virtualization to enable flexible, dynamic bandwidth allocation and increased

headroom while also dramatically reducing the number of physical server adapters needed for a given configuration—which, in turn, helps them to reduce capital and support costs. Ultimately, 10GbE connectivity enables organizations to obtain maximum value from their virtualized server environment. [u](#)

Srinivas Thodati is a senior product marketing manager at Dell and has more than 16 years of experience in the IT industry working for Dell, Hewlett-Packard, Compaq, USRobotics, and 3Com. He specializes in servers and networking for data centers. Srinivas has an M.S. in Electrical Engineering and Computer Science from the University of Illinois as well as a B.S. in Electronics and Communications Engineering.

Brian Johnson is a product marketing engineer at Intel for 10 Gigabit Ethernet products and virtualization technologies, and has more than 12 years of experience in server product planning and marketing. Brian has B.S. degrees in Marketing and Advertising Management from Portland State University; holds CompTIA A+, Server+, and Network+ certifications; and is a Friesen, Kaye, and Associates (FKA) Certified Master Trainer.

MORE
ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Intel Ethernet server adapters:
www.intel.com/go/ethernet
www.intelethernet-dell.com

Intel Virtualization Technology:
www.intel.com/go/vtc

VMware virtualization:
www.vmware.com/virtualization

Dell PowerEdge servers:
dell.com/poweredge

³For more information on how 10GbE connectivity handles the demands of bandwidth-intensive applications, see "New Technologies Speed the Move to 10 Gigabit Ethernet Data Center Connectivity," by Sunil Ahluwalia, in *Dell Power Solutions*, June 2009, dell.com/downloads/global/power/ps2q09-20090230-Intel.pdf.